New DiFX Software Correlator Cluster at Bonn and Summary of Recent Activities

Arno Müskens[1], Walter Alef[2], Alessandra Bertarini[1,2], Simone Bernhart[1], Laura La Porta[1], Helge Rottmann[2], Gabriele Bruni[2], Torben Schüler[3], Walter Brisken[4]

Abstract We report on the status of the Bonn Correlator Center for the period 2015/2016. The correlator center has been operated jointly by the Max Planck Institute for Radio Astronomy (MPIfR) in Bonn and the Federal Agency for Cartography and Geodesy in Frankfurt with support from the Institute of Geodesy and Geoinformation in Bonn for more than 25 years. Since 2010 we have been using the DiFX software correlator [Deller et al. (2011)] for astronomical and geodetic correlation on a High Performance computing Cluster (HPC). We summarize our first experience with correlating on our new upgraded cluster, which was designed to be compatible with the VGOS requirements. The new cluster was purchased in December 2015. We also report on newly implemented features, like playback of Mark 6 recordings, and on other software improvements.

Keywords VLBI correlation, DiFX correlator, geodetic VLBI, VGOS

1 Introduction

Since 1978, the MPIfR has been hosting five generations of VLBI correlators: Mark II, Mark III, Mark IIIA, Mark IV [Whitney et al.(2004)], and the Distributed FX software correlator (DiFX). The correlator center has been operated jointly by the three partners MPIfR, BKG, and IGG.

In 2009 the correlation was switched to DiFX, after 30 years of operating hardware VLBI correlators. The reasons were the significantly lower investment costs and the far greater flexibility of a software-based solution. The DiFX correlator is Open Source and is maintained by a community of developers mostly located at National Radio Astronomy Observatory (NRAO), Netherlands Institute for Radio Astronomy (ASTRON), MPIfR, Commonwealth Scientific and Industrial Research Organisation (CSIRO), and Massachusetts Institute of Technology (MIT) Haystack Observatory. Since 2010, the geodetic correlation has been done exclusively with the DiFX.

2 Background and Present Correlator Status

The DiFX software correlator was developed at Swinburne University in Melbourne, Australia by Adam Deller and collaborators [Deller et al. (2011)]. It has been adapted to the VLBA operational environment by Walter Brisken and NRAO staff, and has been further developed for years by the worldwide DiFX developers group.

DiFX is executed on an HPC, which was upgraded in December 2015 after about seven years of operation. The upgrade was financed by the Max Planck Society with a significant contribution from the German mapping authority BKG. The upgrade went smoothly with a correlator downtime according to plan of less than two weeks. The new cluster is roughly ten times faster than the old one, and will be able to handle wide-band

1. Rheinische Friedrich-Wilhelms Universität Bonn, IGG
2. Max-Planck-Institut für Radioastronomie (MPIfR)
3. Bundesamt für Kartographie und Geodäsie (BKG)
4. National Radio Astronomy Observatory (NRAO)

VLBI observations in the mm-domain and geodesy for at least the next five years.

Some of the features of the new VLBI cluster are:

- 68 compute nodes;
- each node with two processors with 20 compute cores, which adds to 1,360 compute cores in total;
- ∼960 TB disk storage (14 RAIDS; 740 TB for VLBI);
- 56 Gbps Infiniband fabric as interconnect (plus 1 GE Ethernet for service tasks);
- 48 TB disk space for correlated data;
- 40 TB backup for correlated data;
- two head nodes for correlation (several correlations can be run in parallel);
- one head node for other tasks;
- one appliance computer for cluster installation and monitoring;
- 17 Mark 5 units for playback of data into the correlator (all Mark 5 formats);
- six Mark 6 units plus two expansion chassis for 1 mm VLBI data and VGOS data;
- 1 Gbps connection to the Internet reserved for VLBI data transfers (soon 2 Gbps).

In addition to the standard DiFX versions (e.g., latest stable = 2.4, development), which are used for "normal" correlation, a DiFX branch for the correlation of RadioAstron data was developed at MPIfR: DiFX-RA. The most significant part of this effort involved modifying the delay model server (Calc, from the Calc/Solve package) to be able to calculate delay information for telescopes with arbitrary coordinates and velocities (that is, not fixed on the ground), and changes to the DiFX metadata system to deal with the changing position and velocity of the spacecraft as a function of time. The delay model was also modified to correct for (general) relativistic effects as the highly elliptical orbit of the spacecraft results in large changes in velocity and gravitational potential compared to the terrestrial frame. DiFX-RA was used extensively to correlate a number of RadioAstron projects in Bonn, but also to help debug the Astro Space Center (ASC) correlator in Moscow.

To perform the fringe search for RadioAstron data, we installed the PIMA[1] software with the help of L. Petrov and the ASC in Moscow.

[1] http://astrogeo.org/pima

Other enhancements are the implementation of a database for experiment status and disks, and archiving of the raw correlated data together with their FITS-IDI[2] or HOPS[3] exports on the new MPIfR archive server.

We have implemented native Mark 6 playback into the correlator, the handling of the Mark 6 modules by the correlator, and in addition special VDIF modes for the DBBC3 VLBI backend, all of which are available in the latest development release.

At present we are running Mark 6 recorders in native playback mode for data from VGOS and the Event Horizon Telescope (EHT). The usage of multiple data streams per station was implemented by W. Brisken financed by MPIfR. For instance, an EHT observation at 64 Gbps would deliver the data spread over four Mark 6 recorders with extension chassis, resulting in 16 disk modules to be played into the correlator. We expect that DiFX 2.5 with full Mark 6 playback support will become available in a couple of months. Table 1 lists some important capabilities of the Bonn DiFX correlator.

Table 1 Capabilities of the Bonn DiFX correlator.

Feature	Description
Geometric model	Calc 9 (Calc 11 as plugin)
Phase Cal	Phase-cal extraction of all tones in all subbands simultaneously
Pre-averaging time	From milliseconds to seconds
Spectral resolution	Up to 256,000 channels have been used
Signal	Single- and dual-frequency, all four Stokes parameters
Export	Fits export, Interface to MK IV data format for HOPS
Pulsars	Pulsar correlation with incoherent dedispersion

In general the Bonn Correlator will be powerful enough for the requirements of VGOS for the next five years, in particular as the present estimates of the readiness of VGOS antennas seem to be a bit optimistic. We expect that the EHT will deliver 64 Gbps in the next two years, which should not be an issue either.

[2] ftp://ftp.aoc.nrao.edu/pub/software/aips/TEXT/PUBL/AIPSMEMO102.PS

[3] http://www.haystack.mit.edu/tech/vlbi/hops.html

Fig. 1 The correlator room at MPIfR—cluster side. The four racks of the cluster are visible behind the glass wall.

Fig. 2 The correlator room at MPIfR—playback side. The three racks on the left house Mark 5 units. The two racks on the right contain Mark 6 units. A glass wall with door shields the operators from the noise.

3 Correlator Usage 2015

The Bonn correlator is the only VLBI correlator worldwide which is shared by astronomers and geodesists with a usage fraction of about 50% for geodesy and 50% for astronomy.

The geodetic throughput of the Bonn correlator in 2015 was roughly identical to the previous years:

- 52 R1 sessions,
- six EURO sessions,
- seven T2 (with up to 20 antennas) sessions,
- six OHIG sessions,
- four R&D sessions,
- 48 INT3 (in e-VLBI mode) Intensives,
- 30 INT2b (30 days, baseline Ts–Wn; e-VLBI mode) Intensives, and
- several additional correlations for DBBC testing (Onsala, Yebes, Wettzell).

Most of the tests for the DBBC development and performance evaluation are correlated in Bonn. Testing of the DBBCs usually consists of parallel observations (whenever possible) of a geodetic session with both analog and digital backends. If requested, we support every station when new software or hardware for recording or transferring data are introduced like Flexbuff, VDIF format, or jive5ab.

It should be noted that most geodetic observatories do not send disk modules anymore, but transfer the data via Internet.

One part of the astronomy load of the correlator is data from the Global mm VLBI Array (GMVA), which observes two sessions per year of up to five days duration at 3-mm wavelength. Up to 15 antennas participate in GMVA observations. For the last two years the datarate has been increased to 2 Gbps which results in about 500 TB of total disk space recorded in each session.

Correlation of RadioAstron observations has been a big load in the last few years. It requires a very time consuming additional correlation pass, because the RadioAstron clock has to be searched for every scan. As the orbit of RadioAstron is only known to about 500 m, the changes in delay and delay-rate can be quite large, so that an additional acceleration term has to be taken into account. So far, nine observations with up to 30 antennas at 256 Mbps, full track on the source, have been correlated. At times a third correlation pass with an improved satellite orbit was necessary.

Currently the details of correlating high-bandwidth data recorded at 32 Gbps by EHT at 1-mm wavelength are being investigated.

4 Staff Management

The Bonn geodesists have operated their correlation independently of the astronomers from the very beginning. Experiment preparation and post-processing are handled separately for geodetic and astronomical observations.

- **Geodesy:** A. Bertarini, L. La Porta, S. Bernhart, and A. Müskens. Their tasks are:
  - Schedule preparation of nearly all experiments correlated in Bonn.
  - Preparation of correlator control files.
  - Fringe search.
  - Supervision of correlation.
  - Verification of the correlation results and setting up of possible re-correlation.
  - Management of recording media logistics and electronic data transfer.
- **Astronomy:** W. Alef, A. Bertarini, H. Rottmann, G. Bruni, H. Sturm, and R. Märtens.
  - Management of the correlator center.
  - Set-up and verification of astronomical correlation.
  - Cluster management, including operating system and IT security.
  - Maintenance of Mark 5 and Mark 6 playback units as well as of disk modules and cluster hardware.

General computing services and the archive server are provided by the MPIfR computer division.

As a great benefit it turned out that three of the four main geodetic VLBI tasks (*Schedule preparation, Recording, Correlation and Analysis*) are done by our group—the Bonn Geodesists. This allows a tight feedback loop and has resulted in a highly motivated group of people, complemented by the excellent work done by the personnel at the participating geodetic stations.

5 Experiment Distribution among IVS Correlators

Table 2 gives a short overview of the experiment distribution among the IVS correlators worldwide. Unfortunately CRTN (Curtain University software correlator, Australia) is not operational anymore since 2015. Due to the fact that Bonn and WACO have reached their full workload, other IVS correlators need to takeover those experiments. Nevertheless, in the next years several other correlators could become available for IVS correlation. IAA in St. Petersburg/Russia and SHAO (Shanghai VLBI Correlator) plan to upgrade their existing correlator for VGOS correlation. In addition Vienna/Austria plans to acquire a cluster for correlation.

Table 2 Typical workload of the IVS correlators. Note that the Curtin correlator (CRTN) was shutdown in 2015.

Correlator	Percentage	Session type
BONN	32	R1, T2, Eur, OHG, R&D, INT3
WACO	30	R4, CRD, INT1
CRTN	22	AUS-GEO, AUST, AUS-AST
GSI	10	JADE, AOV, JAXA, INT2
SHAO	3	AOV, APSG, CRF
HAYS	2	R&D, T2
NGII	1	AOV

A different common strategy will be needed and a better distribution of the increasing VGOS correlations is required to ensure that the future correlation demand for VGOS can be satisfied. New ideas and discussions about load balancing are essential for the future of the IVS community and the success of the VGOS program. Discussions are in progress whether it is better to have several effective and very well networked correlators spread over the world or only one big correlator center. Both options need a great and well-connected data center (cloud) where all recorded IVS data can be stored and prepared for further processing. It goes without saying that an experienced staff is required to guarantee an excellent scientific correlation and post-correlation analysis for all IVS sessions.

6 e-transfer

Nowadays 90% of the stations e-transfer their data to Bonn. The average amount of e-transferred data per week for the regular R1 and INT3 observations alone is more than 10 TB.

Most transfers utilize the UDP-based Tsunami protocol and the jive5ab (m5copy) script. The achieved data rates range from 100 to 800 Mbps. The present Internet connection to the MPIfR is a 1-Gbps dedicated line to the GÉANT node in Frankfurt. The upgrade to 10 Gbps Internet connection to meet the requirements of VGOS has not yet been realized due to cost issues.

For future VGOS observations we should expect a typical data volume of around 40 TB/day/antenna recorded with 8 Gbps data rate. Therefore the increase of the Internet connection for correlators and stations is essential. In addition, large data buffers have to be provisioned at the stations and at the correlators for each station. In a transition phase Mark 6 modules can be shipped to ease the financial burden of the above two requirements of fast data lines and local large disk storage. But this will lead to a large investment in Mark 6 disk modules and a significant increase in shipping costs. It should be noted that Mark 6 modules can also be used as local storage as implemented, for instance, at Effelsberg.

7 Brain Drain and Outlook

A problem, which is easily overlooked, is that several important VLBI experts will retire in the near future. This is a general problem: VLBI needs new well educated practitioners. The IVS training school is an excellent facility for all young and new members, but experts need many years to become skilled and to gain experience—the special VLBI know-how—which is required.

Some new antennas will become available for use in the next few years. For example, at Wettzell in Germany the new Twin Telescope Wettzell (TTW) for VGOS is close to completion; South Korea has a new antenna for Geodesy at Sejong (NGII, National Geographic Information Institute); in Spain and Portugal, the RAEGE (Atlantic Network of Geodynamical and Space Stations) project aims to establish a network of four fundamental geodetic stations that will fulfill the VGOS specifications. In Spitsbergen/Norway (NMA, Norwegian Mapping Authority) and in Sweden, Onsala Space Observatory, twin telescopes are in the construction phase. Therefore, broadband geodetic sessions at more regular intervals will be scheduled and processed in the near future. The Bonn Correlator will be powerful enough for the requirements of VGOS for the next five years.

References

[Whitney et al.(2004)] A.R. Whitney, R. Cappallo, W. Aldrich, B. Anderson, A. Bos, J. Casse, J. Goodman, S. Parsley, S. Pogrebenko, R. Schilizzi & D. Smythe. Mark 4 VLBI correlator: Architecture and algorithms Radio Science, 2004, 39, 1007.

[Deller et al. (2011)] A. T. Deller, W. F. Brisken, C. J. Phillips, J. Morgan, W. Alef, R. Cappallo, E. Middelberg, J. D. Romney, H. Rottmann, S. J. Tingay & R. Wayth. DiFX-2: A More Flexible, Efficient, Robust, and Powerful Software Correlator, PASP, 2011, 123, 275–287.

[Kardashev et al. (2012)] N. S. Kardashev, Y. Y. Kovalev, & K. I. Kellermann. The Radio Science Bulletin, 2012, 343, 22–29.

[Bruni et al. (2014)] G. Bruni, J. Anderson, W. Alef, A. Lobanov, & A. J. Zensus. Space-VLBI with RadioAstron: new correlator capabilities at MPIfR, in Proceedings of the 12th European VLBI Network Symposium and Users Meeting, Cagliari, Italy, 2014, ID: 119.